

Dynamic Windows Scheduling with Reallocation [★]

Martín Farach-Colton¹, Katia Leal², Miguel A. Mosteiro³, and Christopher Thraves²

¹ Dept. of Computer Science, Rutgers University, Piscataway, NJ, USA & Tokutek Inc.,
farach@cs.rutgers.edu

² GSyc, Universidad Rey Juan Carlos, Madrid, Spain,
{katia, cbthraves}@gsyc.es

³ Dept. of Computer Science, Kean University, Union, NJ, USA,
mmosteir@kean.edu

Abstract. We consider the Windows Scheduling problem. The problem is a restricted version of Unit-Fractions Bin Packing, and it is also called Inventory Replenishment in the context of Supply Chain. In brief, the problem is to schedule the use of communication channels to clients. Each client c_i is characterized by an active cycle and a window w_i . During the period of time that any given client c_i is active, there must be at least one transmission from c_i scheduled in any w_i consecutive time slots, but at most one transmission can be carried out in each channel per time slot. The goal is to minimize the number of channels used.

We extend previous online models, where decisions are permanent, assuming that clients may be reallocated at some cost. We assume that such cost is a constant amount paid per reallocation. That is, we aim to minimize also the number of reallocations. We present three online reallocation algorithms for Windows Scheduling. We evaluate experimentally these protocols showing that, in practice, all three achieve constant amortized reallocations with close to optimal channel usage. Our simulations also expose interesting trade-offs between reallocations and channel usage. We introduce a new objective function for WS with reallocations, that can be also applied to models where reallocations are not possible. We analyze this metric for one of the algorithms which, to the best of our knowledge, is the first online WS protocol with theoretical guarantees that applies to scenarios where clients may leave and the analysis is against current load rather than peak load. Using previous results, we also observe bounds on channel usage for one of the algorithms.

Keywords: Reallocation Algorithms, Windows Scheduling, Radio Networks, Unit Fractions Bin Packing

[★] This work was supported in part by the National Science Foundation (CCF- 0937829, CCF- 1114930) and Kean University UFRI grant.

1 Introduction

We study the Windows Scheduling⁴ problem (WS) [4, 5, 8], which is a restricted version of Unit-Fractions Bin Packing (UFBP) [1, 2, 5, 7, 9, 11], and it is also called Inventory Replenishment in the context of Supply Chain [17]. In brief, the WS problem is to schedule the use of communication channels to clients. Each client c_i is characterized by an active cycle (with arrival and departure times) and a window w_i also called laxity⁵ (in the context of job scheduling). During the period of time that any given client c_i is active, there must be at least one transmission from c_i scheduled in any w_i consecutive time slots. The optimization criterion is to minimize the number of channels used. The WS problem appears in many areas such as Operations Research, Networks, Streaming, etc. More application details can be found in [4, 5, 8, 17] and the references therein.

Given that even a restricted version of the WS problem was shown to be NP-hard even for one channel [3], practical WS solutions are only approximations. In the WS literature, competitive analysis of the ratio between the number of channels used by an online algorithm with respect to an optimal algorithm has been carried out in two flavors: the maximum ratio for any given time instant (also called *against current load*), and as the ratio of the maxima (also called *against peak load*). In the model of [4, 5] both competitive ratios are the same because clients do not leave the system. In [8, 17] the competitive ratio is against peak load. In the present work, we carry out competitive analysis against current load.

In [4, 5] clients do not leave the system whereas in [8, 17] clients may leave. We further extend the model assuming that clients may be reallocated at some cost. As used in [6] (for Job Scheduling), we call this class of protocols **Reallocation Algorithms**. Reallocation algorithms are a middle ground between online algorithms (where assignments are final) and offline algorithms which can be used repeatedly if reallocations are free. Reallocation has been studied previously for Load Balancing [13, 16], and for UFBP in [1, 2] where it was called *semi-online algorithms*. In [11], also for UFBP, the reallocation cost paid is only computation time. Preemptive Job Scheduling has also been studied assuming that a cost for preemption is paid [10, 12, 14] but preempted jobs in this model are left idling. In the present work, we assume that the cost of reallocation is a constant amount. That is, we aim to minimize the number of reallocations.

In this paper, we present three protocols for online WS with reallocation. Namely, Preemptive Reallocation, Lazy Reallocation, and Classified Reallocation. In brief, the first is a repeated (upon each client arrival or departure) application of the offline protocol in [4]. Aiming to minimize channel usage, clients are preemptively reallocated to guarantee the same offline packing. Instead, in Lazy Reallocation, clients are not reallocated as long as a maximum number of channels in use is not exceeded. The idea is

⁴ Were it not for the extensive literature on windows scheduling, we would have chosen a different name. After all, the goal of windows scheduling is not to schedule windows but transmissions. Moreover, those transmissions need not be scheduled within fixed epochs, as the term *window* seems to suggest. Nonetheless, we use the current notation for literature consistency.

⁵ Throughout the paper, we will use the term *laxity* instead of *window* since it conveys better the concept of maximum delay between transmissions. Nonetheless, we will also use the term *window* when describing the related literature.

to save reallocations taking advantage of all possible channels. Finally, Classified Re-allocation is designed to guarantee an amortized constant number of reallocations. The main approach is to classify clients by laxity.

We evaluate experimentally all three protocols. Our simulations show that, in practice, all three achieve constant amortized reallocations with close to optimal channel usage. Our simulations also expose interesting trade-offs between reallocations and channel usage. On the theoretical side, we introduce a new objective function for WS with reallocations, that can be also applied to models where reallocations are not possible. We analyze this metric for Classified Re-allocation which, to the best of our knowledge, is the first online WS protocol for dynamic scenarios (clients may leave) with theoretical guarantees (against current load). Using previous results, we also observe bounds on channel usage for Preemptive Re-allocation.

The rest of the paper is organized as follows. First we overview the related work in Section 2. Section 3 includes formal definitions of the problem and the model. The details of our contributions are presented in Section 4. We present and analyze reallocation protocols in Sections 5 and 6, and simulations in Section 7.

2 Related Work

WS is a problem with applications to various areas such as communication networks, supply chain, job scheduling, media on demand systems, etc. We overview here the most related work.

In [4] the authors studied WS for broadcast systems where pages have to be allocated to slotted broadcast channels. The allocation comprises a schedule of pages to slots such that the gap between two consecutive slots reserved for page i is at most its window w_i . The model is static in the sense that, once pages arrive to the system, they do not leave. Two optimization criteria are studied: to minimize the number of channels used by a fixed set of clients and to maximize the number of clients allocated to a fixed set of channels. For the former problem they show an offline approximation of $H + O(\ln H)$ where $H = \sum_i 1/w_i$. The approach followed is to round down each window size to the largest number of the form $x2^y$, where x is an odd positive integer and y is a non-negative integer. Then, allocate all clients of the same window size to the same channel creating a new one whenever necessary. For online algorithms, an approximation of $H + O(\sqrt{H})$ channels was shown later in [5]. The approach is similar but the analysis takes into account that, being online, pages are allocated one by one and decisions are final.

A scenario where pages may leave the system was termed WS with Temporary Items in [8]. Bounds in the latter work were proved on the competitive ratio against peak load. That is, the online and the offline channel-usage maxima compared may occur at different times. That is, taking advantage of (rather than overcoming) windows departures. Specifically, the authors showed that any online algorithm has a competitive ratio at least $1 + 1/w_{\min} - \epsilon$, for any $\epsilon > 0$, independently of knowledge of departure time. On the positive results side, they show an online algorithm that is $2 + 2/(\lfloor w_{\min} \rfloor - 1)$ -competitive, for $w_{\min} \geq 2$. Recently [17], this upper bound was improved to $4w_{\min}/(3w_{\min} - 4)$, for $w_{\min} \geq 4$, in the context of supply chain

inventory replenishment. The algorithm makes use of a combined first fit policy for assignment.

In [5], WS was studied as a *restricted* version of UFBP. The reason is apparent, the combination of windows in WS is not additive as in UFBP. Hence, UFBP upper bounds do not apply to WS but lower bounds do. A lower bound of $H + \Omega(\ln H)$ on the online approximation of UFBP, which then applies also to WS, was shown in [5]. A dynamic version of UFBP where items leave was studied in [7]. The authors show an upper bound of 3 on the competitive ratio of best fit and worst fit allocation policies, and an upper bound of 2.4942 for first fit. UFBP with reallocation was studied in [1, 2] as *semi-online* algorithms with *repacking*. In their model items do not leave the system, hence only lower bounds apply to WS with reallocation, but only upper bounds are presented. In [11], *fully dynamic* UFBP algorithms apply to settings where items may leave, but the reallocation cost is only bounded in terms of computation time.

3 Model

In this Section, we describe the Windows Scheduling problem and introduce the notation used throughout the paper.

We assume that time is discrete and it is determined by a global clock that runs in *time steps*. Consider a set of clients \mathcal{C} that require periodical radio transmissions. Each client $c_i \in \mathcal{C}$ is determined by three parameters: **arrival** time t_i , **departure** time d_i , and **laxity** w_i . We say that client c_i is **active** during the time interval $[t_i, d_i]$. We assume that there is an infinite number of available radio channels $\mathcal{R} = \{R_j\}_{j \in \mathbb{N}}$. During one time step, only one client is able to transmit via a single channel.

A **schedule** for a client $c_i \in \mathcal{C}$ is a set $S(c_i) \triangleq \{(R_k^i, t_k^i)\}_{k \in [1, \ell_i]}$, where, channel R_k^i is the reserved channel for c_i 's transmission at time step t_k^i . Time steps t_k^i should satisfy the following conditions:

$$t_1^i - t_i \leq w_i, \quad (1)$$

$$d_i - t_{\ell_i}^i \leq w_i \text{ and} \quad (2)$$

$$t_{k+1}^i - t_k^i \leq w_i \text{ for all } 1 \leq k < \ell_i. \quad (3)$$

Parameter ℓ_i denotes the number of transmissions performed by client c_i . Note that this parameter is independent for each client, it is not set before hand and has to be set by the schedule so that conditions (1), (2) and (3) are satisfied. Thus, conditions (1), (2) and (3) ensure that each client c_i transmits only while it is active and that periodical transmissions are so that the laxity condition w_i for client c_i is satisfied. Since client c_i has to transmit periodically every (at most) w_i time steps, we say that the **load of client** c_i is $1/w_i$.

We say that schedule $S(c_i)$ **reallocates** client c_i each time that $R_k^i \neq R_{k+1}^i$ for $1 \leq k < \ell_i$. We denote by $reall(S(c_i))$ the number of times schedule $S(c_i)$ reallocates client c_i . The total **number of reallocations** of schedule $S(\mathcal{C})$ for the set of clients \mathcal{C} is denoted by $reall(S(\mathcal{C}))$ and is defined as follows:

$$reall(S(\mathcal{C})) \triangleq \sum_{c_i \in \mathcal{C}} reall(S(c_i)).$$

We are interested in the *amortized* number of reallocations, that is, the number of reallocations of a schedule divided by the number of arrivals and departures in the system. For that purpose, let us define a **round** as the set of time slots between events (arrival or departure of a client). In other words, at each round occurs exactly one event, either an arrival or a departure of a client. Hence, if $reall_r(S(\mathcal{C}))$ denotes the number of reallocations up to round r produced by schedule $S(\mathcal{C})$, the *amortized* number of reallocations of schedule $S(\mathcal{C})$ in the set of clients \mathcal{C} at round r is $reall_r(S(\mathcal{C}))/r$.

Naturally, there exist a trade-off between the number of reallocations and the number of channels used by a schedule, the more reallocations the less number of channels used at each round. Indeed, if a schedule can reallocate clients freely, it can achieve the optimal offline number of channels used at each round by simply computing the optimal offline schedule at each round, and reorganizing the schedule via reallocations in order to mimic the optimal offline schedule. Thus, we are interested in understanding this trade-off at each round. For any round r , let $\mathcal{C}(r) \subseteq \mathcal{C}$ be the set of active clients in the system at round r . Hence, we define $H_r \triangleq \sum_{c_i \in \mathcal{C}(r)} 1/w_i$ as the **load of the system** at round r . The value $\lceil H_r \rceil$ is a lower bound on the optimal number of used channels at round r . Hence, $\mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil$ is an upper bound on the competitive ratio of the number of channels used at round r , where $\mathcal{R}(S(\mathcal{C}))_r$ denotes the number of channels used at round r by schedule $S(\mathcal{C})$.

Using this notation, we define the online WS problem with reallocations as follows:

Definition 1. Let \mathcal{C} be a set of clients revealed online, that is, arrivals and departures of clients are revealed one by one. The **Online Windows Scheduling with Reallocations** problem is to determine a schedule $S(\mathcal{C})$, possibly with reallocations, so that the following sum is minimum.

$$\frac{reall_r(S(\mathcal{C}))}{r} + \frac{\mathcal{R}(S(\mathcal{C}))_r}{\lceil H_r \rceil}$$

Notice that this metric is against current load and applies to any given round r . Should we instead be interested in the maximum competitive ratio, we could simply plug the maximum up to round r instead of the current. Likewise, should the reallocation cost be any known value c , it could be simply introduced in the expression above multiplying the first term.

4 Our Contributions

The main contributions of this paper follow. It should be noticed that our upper bounds cannot be compared with previous theoretical work on WS because either their models assume that clients do not leave [4,5], or the bounds on channel usage are proved against peak load [8, 17], or they apply to the less restrictive UFBP problem [1, 2, 5, 7, 11].

- The presentation of three protocols for Online Windows Scheduling with Reallocation called Preemptive Reallocation, Lazy Reallocation, and Classified Reallocation. Preemptive Reallocation guarantees low channel-usage because clients

are preemptively reallocated to achieve the offline packing of [4]. Attempting to save reallocations while keeping the channel usage bounded, in Lazy Reallocation clients are not reallocated as long as a maximum number of channels in use is not exceeded. In Classified Reallocation, the main approach is to classify clients by laxity, except for “large” laxities that are allocated in one special channel to maintain the channel-usage overhead below an additive logarithmic factor. To the best of our knowledge, Classified Reallocation is the first online WS protocol for dynamic scenarios (clients may leave) with theoretical guarantees (against current load).

- The experimental evaluation of all three protocols showing that, in practice, all of them achieve constant amortized reallocations with close to optimal channel usage. Our simulations also expose interesting trade-offs between reallocations and channel usage.
- The introduction of a new objective function for Online Windows Scheduling with Reallocations, that can be also applied to models where reallocations are not possible. This metric combines linearly the effect of reallocations amortized over rounds with the number of channels used in contrast with the optimal UFBP, which is a lower bound on the WS optimal allocation.
- Using previous results [4], for Preemptive Reallocation we observe an upper bound of $\lceil 2H_r \rceil$ on channel usage for every round r . Given that in Lazy Reallocation the maximum number of channels is a parameter, its channel usage depends on the implementation.
- We prove that, for any round r , the number of reallocations required by Classified Reallocation is at most $3r/2$, and the number of channels used is at most $OPT(r) + 1 + \log(\min\{w_{\max}(r), \lceil \lceil n(r) \rceil \rceil\} / w_{\min}(r))$ ⁶, where $OPT(r)$ is the optimal number of channels required, $w_{\max}(r)$ and $w_{\min}(r)$ are respectively the maximum and minimum laxities in the system, and $n(r)$ is the number of clients in the system, all for round r . We apply these bounds to our objective function in Definition 1.

5 Preemptive Reallocation and Lazy Reallocation Algorithms

The first two algorithms use the concept of a *broadcast tree* to represent the schedule corresponding to each channel.

Broadcast Tree to Represent a Schedule. Similarly to the the work by Bar-Noy et al. [4] and Chan et al. [8], we represent a schedule for a channel with a tree. In particular, we use a binary tree where all nodes have exactly zero or two children. Each leaf of the binary tree has assigned one client (different for each leaf) or the leaf is empty. Each complete binary tree with such an assignment represents uniquely one schedule for a channel. Given a complete binary tree, the schedule picks one leaf at each time step so that either the assigned client transmits or no transmission is produced if the leaf is empty. In order to pick one leaf, walk down from the root up to one leaf following the next recursive rules: the first time a bifurcation is visited go to the left child. The i -th

⁶ Throughout, \log means \log_2 unless otherwise stated. We define $\lceil \lceil n(r) \rceil \rceil$ as the smallest power of 2 that is not smaller than $n(r)$.

time a bifurcation is visited go to the child that was not visited in the previous visit. Such trees are called **broadcast trees**.

A used channel will be fully described by its broadcast tree in the following two algorithms. Thus, the following two algorithms are determined by (i) the procedure to assign clients to a leaf in a broadcast tree when a client arrives and (ii) the procedure to reallocate clients if required when a client leaves the system.

Greedy Construction of a Broadcast Tree. Consider an empty channel as a single node, the root of its corresponding tree. Define every root as **available**. For the first client c_1 that arrives, let v_1 be the nonnegative integer such that $2^{v_1} \leq w_1 < 2^{v_1+1}$. Append to any root a binary tree of height v_1 in which for each depth from 1 to $v_1 - 1$, there is a single leaf and for depth v_1 there are two leaves. Set every new leaf as available except for one leaf at depth v_1 which is assigned to c_1 . For the i -th client c_i that arrives, let v_i be the nonnegative integer such that $2^{v_i} \leq w_i < 2^{v_i+1}$. If there is an available leaf at depth v_i , assign c_i to that leaf. Otherwise, let $0 < u < v_i$ be any value such that there is an open leaf at depth u in some used broadcast tree. Append a binary tree of height $v_i - u$ to that leaf in which for each depth from u to $v_1 - u - 1$, there is a single leaf and for depth v_1 there are two leaves. Set every new leaf as available except for one leaf at depth v_i which is assigned to c_i . If no such u exists, append the described tree to an available root. In this case consider $u = 0$.

Preemptive Reallocation. The **Preemptive Reallocation** procedure maintains the following invariant: for each depth there is at most one broadcast tree with a leaf available. If the invariant is violated after some departure, it means that there are two broadcast trees with an available leaf in the same depth. Then, the branches that hang from the twin node of the available leaf can be hanged from the same tree to reinstate the invariant. The Preemptive Reallocation procedure does so minimizing the total number of reallocations, that is, moving the branch that hangs from the broadcast tree with less clients assigned.

Lazy Reallocation. The **Lazy Reallocation** procedure reallocates only when the fraction $\mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil$ exceeds a threshold T after one departure. It exhaustively reallocates clients until no more reallocations can be made. Reallocations are done according to the invariant used by the Preemptive Reallocation procedure merging the smallest depth with two available leaves at the same depth.

The **Preemptive Reallocation** and **Lazy Reallocation** algorithms use the greedy construction of broadcast trees when clients arrive and preemptive and lazy reallocations procedures when clients leave the system, respectively. Lemma 5 of Bar-Noy et al. [4] implies that the Preemptive Reallocation algorithm guarantees $\mathcal{R}(S(\mathcal{C}))_r < \lceil 2H_r \rceil$ for every round r . Hence, $\mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil < 2$ for every round r when $S(\mathcal{C})$ is constructed according to the Preemptive Reallocation algorithm. On the other hand, by definition of the Lazy Reallocation procedure, the Lazy Reallocation algorithm guarantees that $\mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil \leq T$ when $S(\mathcal{C})$ is constructed according to the Lazy Reallocation algorithm. In Section 7, we study via experiments the behavior of $reall_r(S(\mathcal{C}))/r$ for the Preemptive Reallocation and Lazy Reallocation algorithms. In the particular implementation of the Lazy Reallocation algorithm included in this work, we set the threshold T equal to $4\sqrt{H_r}$. Hence, in that case it holds $\mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil \leq 4\sqrt{H_r}$.

For both algorithms, Preemptive Reallocation and Lazy Reallocation, we show experimentally that $reall_r(S(\mathcal{C}))/r \leq 1$.

6 Classified Reallocation

In this section, we present a reallocation algorithm that guarantees $O(1)$ reallocations amortized on rounds. The details of the protocol can be found in Algorithm 1. Bounds on channel usage and reallocations are proved in Theorem 1. Corollary 1 establishes these bounds in our objective function. For convenience, for any number x , we define the **hyperceiling** of x , denoted as $\lceil\lceil x \rceil\rceil$, to be the smallest power of 2 that is not smaller than x . For this algorithm, we restrict the input to laxities that are powers of 2. The study of inputs with arbitrary laxities is left for future work.

The intuition of the protocol is the following. When a new client c_i arrives and there are already $n - 1$ clients in the system, for $n \geq 1$, we distinguish two cases. If the laxity of c_i is at least $2\lceil\lceil n \rceil\rceil$, assign c_i to a special channel called **big channel**. All clients allocated to the big channel transmit with period $\lceil\lceil n \rceil\rceil$ so, because there are n clients in the system, one big channel is enough. All the other channels being used are called **small channels**. Otherwise, if the laxity is $w_i < 2\lceil\lceil n \rceil\rceil$, assign client c_i to a channel reserved for laxities w_i , we call it **w_i -channel**. If such channel does not exist or all w_i -channels are full, reserve a new one. For any laxity w_i , all clients allocated to a w_i -channel transmit with period w_i . That is, a maximum of w_i clients can be allocated to a w_i -channel. When a client c_j of laxity w_j leaves a channel C , if C is the big channel do nothing. Otherwise, reallocate a client from the w_j -channel of minimum load (if any other) to the slot left by c_j .

With each arrival or departure the number of clients n change. If, upon an arrival, $\lceil\lceil n \rceil\rceil$ becomes larger than the laxity of some clients allocated to a big channel, reallocate those clients to other channels according to laxity, reserving new channels if necessary. Because n was doubled since the allocation of these clients, these reallocations are amortized by the arrivals that doubled n . If, upon a departure, $2\lceil\lceil n \rceil\rceil$ becomes smaller than the laxity of some clients, reallocate those clients to a big channel, releasing the reservation of the channels that become empty. Because n was halved since the allocation of these clients, these reallocations are amortized by the departures that halved n .

The following theorem bounds the number of reallocations and the number of channels used by the Classified Reallocation algorithm.

Theorem 1. *Given a set of clients \mathcal{C} , the schedule $S(\mathcal{C})$ obtained by the Classified Reallocation algorithm requires at most $3r/2$ reallocations up to round r . Additionally, for any round r such that $\mathcal{C}(r) \neq \emptyset$, the number of reserved channels is at most*

$$OPT(r) + 1 + \log \frac{\min \left\{ \max_{i \in \mathcal{C}(r)} \{w_i, \lceil\lceil \mathcal{C}(r) \rceil\rceil\} \right\}}{\min_{i \in \mathcal{C}(r)} w_i},$$

where $OPT(r)$ is the minimum number of channels required to allocate the clients in $\mathcal{C}(r)$.

Algorithm 1: $O(1)$ reallocations. $\lceil \lceil n \rceil \rceil$ is the largest power of 2 that is not greater than n . For any laxity w , all clients allocated to a w -channel are scheduled to transmit with period w .

```

1  $n \leftarrow 0$  // active clients count
2  $\tau \leftarrow 2$  // big channel threshold
3 start tasks 1 and 2
4 Task 1
5   upon arrival of client  $c_i$  do
6      $n \leftarrow n + 1$ 
7     if  $2\lceil \lceil n \rceil \rceil > \tau$  then // consolidate the big channel
8        $\tau \leftarrow 2\lceil \lceil n \rceil \rceil$ 
9       foreach client  $c_j$  in the big channel such that  $w_j < \tau/2$  do
10         if all  $w_j$ -channels are full then reserve a new  $w_j$ -channel
11         reallocate  $c_j$  to the  $w_j$ -channel of minimum load
12       foreach client  $c_j$  in the big channel do
13         re-schedule  $c_j$  to transmit with period  $\tau/2$ 
14     if  $w_i \geq \tau$  then // allocate new client
15       allocate  $c_i$  to the big channel to transmit with period  $\tau/2$ 
16     else
17       if all  $w_i$ -channels are full then reserve a new  $w_i$ -channel
18       allocate  $c_i$  to the  $w_i$ -channel of minimum load
19
20 Task 2
21   upon departure of client  $c_i$  from channel  $c$  do
22      $n \leftarrow n - 1$ 
23     if  $c$  is not the big channel then // consolidate  $w_i$ -channels
24       if  $c$  is empty then release  $c$ 
25       else if there is a  $w_i$ -channel  $c' \neq c$  that is not full then
26         reallocate a client from  $c'$  to  $c$ 
27     if  $2\lceil \lceil n \rceil \rceil < \tau$  then // consolidate the big channel
28        $\tau \leftarrow 2\lceil \lceil n \rceil \rceil$ 
29       foreach client  $c_j$  in the big channel do
30         re-schedule  $c_j$  to transmit with period  $\tau/2$ 
31       foreach  $w$ -channel  $c''$  such that  $w > 2\tau$  do
32         foreach client  $k$  in  $c''$  do
33           reallocate  $k$  to the big channel to transmit with period  $\tau/2$ 
34         release  $c''$ 

```

Proof. The bound on the number of channels follows from the algorithm. Specifically, for any round r , there are at most $OPT(r)$ w_i -channels full, and there is at most one big channel. With respect to the not-full w_i -channels, the maximum w_i is either $\max_{i \in \mathcal{C}(r)} w_i$ or $\lceil |\mathcal{C}(r)| \rceil$, whatever is smaller. Given that all laxities are powers of 2 the bound follows.

To bound the reallocations, we map reallocations to arrivals or departures. Given that rounds are defined by arrivals and departures, the amortized bound follows. The mapping is the following. Clients are reallocated due to one of three possible events as follows.

1. Some client c_j with laxity w_j departed from a w_j -channel c that was full. Then, if there is some other w_j -channel c' that is not full, some client c_i allocated to c' is reallocated to the slot left by c_j in c (see Lines 25-26 in Algorithm 1).
2. Upon the arrival of some client, the total number of clients n increases making $2\lceil n \rceil$ larger than the big-channel threshold. Then, any client c_i allocated to the big channel whose laxity is $w_i < \lceil n \rceil$ is reallocated to a w_i -channel (see Lines 7-11 in Algorithm 1).
3. Upon the departure of some client, the total number of clients n decreases making $2\lceil n \rceil$ smaller than the big-channel threshold. Then, all clients in all w -channels such that $w > 4\lceil n \rceil$ are reallocated to the big channel (see Lines 31-34 in Algorithm 1).

No other event triggers a reallocation.

Now, we define the mapping. For Event 1, the reallocation of c_i is mapped to the departure of c_j . To define the mapping for Event 2, we need the following lemmas.

Lemma 1. *Consider a client c_i that has to be reallocated from the big channel in Lines 9-11 of Algorithm 1. Let n be the number of clients in the system at the time of reallocation in Line 11, and n' be the number of clients in the system at the time of the last allocation of c_i . Then, it is $n \geq 2n'$. That is, after the last allocation of c_i to the big channel, the total number of clients in the system at least has doubled.*

Proof. In the following, all line numbers refer to Algorithm 1. Clients are reallocated from the big channel because their laxities are strictly smaller than $\tau/2 = \lceil n \rceil$ (see Line 9). In order to be allocated to the big channel, c_i must have a laxity at least $2\lceil n' \rceil$ if it was upon arrival (see Line 15), or at least $4\lceil n' \rceil$ if it was upon consolidation of the big channel (see Line 31). In either case, it must be $2\lceil n' \rceil \leq w_i < \lceil n \rceil$, which implies that $n \geq 2n'$. \square

Lemma 2. *In any given round, at most half of the clients in the system are reallocated from the big channel after executing Lines 9-11 of Algorithm 1.*

Proof. Out of the set of clients being reallocated from the big channel at a given round r , consider the client c_i that was allocated last, say, in some round $r' < r$. From Lemma 1, we know that between r' and r the number of clients in the system has at least doubled. Furthermore, we know that none of the clients that arrived after r' has to be reallocated in round r , because c_i was the last one. Hence, it cannot be that more than half of the clients in the system are reallocated from the big channel in round r . \square

Now we define the mapping for Event 2. Let n be the number of clients in the system at the time of Event 2. As shown in Lemma 2, *at most* $n/2$ clients have to be reallocated. And, as shown in Lemma 1, after the last allocation to the big channel of any client that has to be reallocated, the total number of clients in the system *at least* has doubled. Hence, the at most $n/2$ reallocations are mapped to the at least $n/2$ arrivals.

To define the mapping for Event 3, we need the following lemma.

Lemma 3. *Consider a client c_i that has to be reallocated to the big channel in Lines 31-34 of Algorithm 1. Let n be the number of clients in the system at the time of reallocation in Line 33, and n' be the number of clients in the system at the time of the last allocation of c_i . Then, it is $n \leq n'/2$. That is, after the last allocation of c_i , the total number of clients in the system at least has halved.*

Proof. In the following, all line numbers refer to Algorithm 1. Clients are reallocated to the big channel because their laxities are strictly larger than $2\tau = 4\lceil n \rceil$ (see Line 31). Because c_i was not in the big channel, c_i must have a laxity strictly smaller than $2\lceil n' \rceil$ if the last allocation was upon arrival (see Line 17), or strictly smaller than $\lceil n' \rceil$ if it was upon consolidation of the big channel (see Line 9). In either case, it must be $4\lceil n \rceil < w_i < \lceil n' \rceil$, which implies that $n \leq n'/4 \leq n'/2$. \square

Finally, the mapping for Event 3 is the following. Let n be the number of clients in the system at the time of Event 3. As shown in Lemma 3, after the last allocation of any client that has to be reallocated, the total number of clients in the system at least has halved. Hence, the at most n reallocations are mapped to the at least n departures.

In the mapping above, there is at most one reallocation for each arrival and at most two reallocations for each departure. Given that there cannot be more departures than arrivals, the number of reallocations up to round r are at most $r/2 + r$. Hence, the claim follows. \square

The following corollary is a direct consequence of Theorem 1 and the fact that $OPT(r) \geq \lceil \sum_{i \in \mathcal{C}(r)} 1/w_i \rceil = \lceil H_r \rceil$.

Corollary 1. *Given a set of clients \mathcal{C} , the schedule $S(\mathcal{C})$ obtained by the Classified Re-allocation algorithm achieves, for any round r , is $reall_r(S(\mathcal{C}))/r + \mathcal{R}(S(\mathcal{C}))_r / \lceil H_r \rceil \leq 5/2 + (1 + \log(\min\{\max_{i \in \mathcal{C}(r)} w_i, \lceil |\mathcal{C}(r)| \rceil\} / \min_{i \in \mathcal{C}(r)} w_i)) / \lceil H_r \rceil$.*

7 Simulations

Model. The deployment of the proposed algorithms on a real environment will require involvement of a large number of active users and resources, which is very hard to coordinate and build, and would prevent repeatability of results. Thus, simulation appears to be the easiest way to analyze the different proposed online reallocation strategies. Based on the simulation results, we can later encourage or discourage the deployment on a real production environment. Next, we present the simulation model implemented for evaluating the performance of the previously proposed reallocation policies, named *CommunicationChannelsSim* (CCSim). The simulated communication channels have

been performed by means of `SimJava 2.0` [15]. `SimJava` is a discrete event, process oriented, simulation package. It is an API that augments Java with building blocks for defining and running simulations.

All the simulations are performed from the point of view of one user. This user submits *clients* to the scheduler. The clients are loaded from an XML input file. For our experiments we have created three different input files each containing 4000 clients with different laxities. These scenarios represent common situations on which clients with different laxities arrive to the system and departure in an instant of time that depends directly on the laxity demanded by the client. The better the laxity, the sooner he will leave the system. In the same way, the scheduler performs one of the proposed re-allocation strategies and acts accordingly when a client arrives or leaves the system. We have implemented three versions of CCSim that only differ in the reallocation policy implemented. As a result, we isolate the reallocation strategy as the only factor that can cause number of channels and number of reallocations variations between these three CCSim versions.

As we mentioned before, we have created our own input files. We have defined the structure of the files using the XML language. Thus, each input file contains 4000 clients with the following characteristics:

- * `id`: each client has its own identifier to differentiate it from the rest.
- * `t_arrive`: arrival simulation time of the client in seconds, for example, 899 seconds.
- * `size`: laxity of the client as a real value, for example, 6.628461669685978. The laxity can be calculated using a Gaussian (normally) distributed double value or it can be chosen uniformly. Thus, we can generate input files following Normal, Uniform or mixed laxity distributions.
- * `w_size`: the laxity is later rounded down to the larger power of 2, for our example, its value would be 4.
- * `t_leave`: simulation time instant in seconds at which the client leaves the system. This value is calculated based on the laxity. Thus, if the client has a laxity less than or equal to 30, `t_leave` will be `t_arrive` plus a random number generated uniformly between 500 and 1000. Conversely, if the client has a laxity greater than 30, `t_leave` will be `t_arrive` plus a random number generated uniformly between 1000 and 1500. As a consequence, clients with a smaller laxity will remain less time in the system. For instance, the client of the example will leave the system at second 1737.0.

Explanations for the main CCSim participating entities and its simulation setup follow.

CCSim represents the complete simulation, and is responsible for the creation of the main simulated entities: Scheduler and User. When the simulation starts, CCSim creates 1 User and 1 Scheduler.

The **Scheduler** entity represents a generic scheduler implementing the corresponding reallocation algorithm. When a Client *arrives*, the Scheduler allocates it in a Broadcast Tree Node. This operation may involve creating a new Tree. Also, when a Client leaves, the Scheduler removes it from the corresponding Tree Node. This operation may imply reallocating Clients and deleting Trees.

The **User** models a user that submits Clients to the Scheduler. The User is responsible of sending the only two system events: Client arrival and Client departure by using Client's `t_arrive` and `t_leave` times.

The **Client** entity represents a generic Client submitted to the Scheduler. This entity provides specific information about each client as defined in the XML input file.

The Broadcast **Tree** entity, as described in Section 5, represents a channel with up to five different levels representing five Client laxities (2, 4, 8, 16, 32). Thus, at each level there are at most 2^i Nodes, with i being a number between 1 and 5. Clients are allocated in their corresponding Node level according to their rounded down laxity `w_size`. Each tree is characterized by a root Node.

The **Node** entity represents a Tree Node in which a Client can be allocated. Each node knows its level or laxity, if it is a root node, a left node or a right node. Also, each node points to its left subtree and to its right subtree.

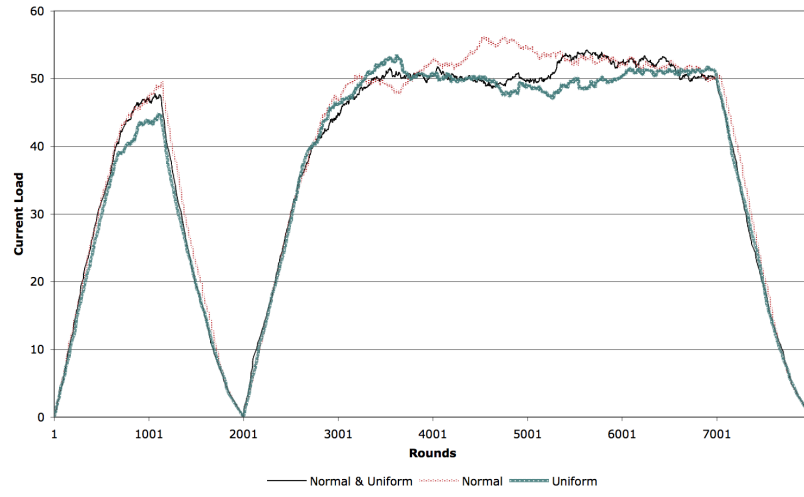
Discussion. Our simulations show similar behaviors for normal and uniform input distributions. All three protocols achieve constant amortized reallocations with close to optimal channel usage. Figures 1 and 2 illustrate the performance of the algorithms along rounds for a normally distributed input. Figure 3 illustrates the trade-offs between channel usage and reallocations for uniform and normally distributed laxities.

We can observe in Figure 1(a) the overall load of the system for a representative input. The load was increased over 1000 rounds until reaching a peak and decreased during the following 1000 rounds until reaching low load. Afterwards, the increase/decrease procedure was repeated but now maintaining the system loaded for approximately 4000 rounds. We can see in Figure 2(b) that Preemptive Reallocation incurs in more reallocations than Lazy Reallocation and Classified Reallocation, but still for all three the amortized reallocations are below 1. When the system is loaded, we observe in Figure 2(a) that the competitive ratio stays around 1.5 for all algorithms. When the system has low load, Preemptive Reallocation still maintains the competitive ratio below 2, whereas Lazy Reallocation increases up to the parametric maximum channel usage. For these simulations, that maximum was $H_r + 4\sqrt{H_r}$, that is, the online maximum channel usage of [5] for clients that do not leave. With respect to Classified Reallocation with low load, the competitive ratio is higher because, for simplicity, the simulation was carried out without using a big channel (refer to the Classified Reallocation algorithm). Preemptive Reallocation and Lazy Reallocation reflect clearly the trade-off between reallocations and channel usage. While the former incurs in more reallocations than the latter, the reverse is true with respect to channel usage. This trade-off becomes more dramatic if the maximum number of channels allowed in Lazy Reallocation is increased. Classified Reallocation, on the other hand, which out of the three is the only algorithm providing theoretical guarantees, has low amortized reallocations (less than 0.5) while maintaining a low competitive ratio when the system is loaded. For systems with frequent low load, the big channel should be implemented. Figure 1(b) illustrates the combination of these factors in our objective function.

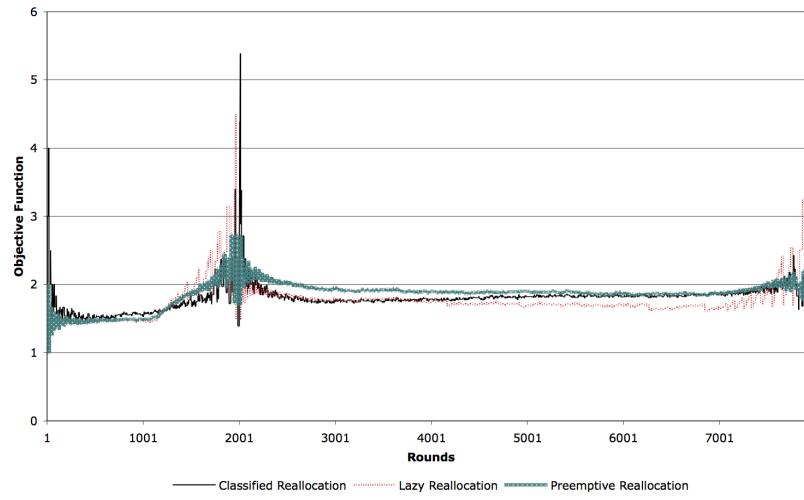
References

1. János Balogh and József Békési. Semi-on-line bin packing: a short overview and a new lower bound. *Central European Journal of Operations Research*, pages 1–14, 2012.

2. János Balogh, József Békési, Gábor Galambos, and Gerhard Reinelt. On-line bin packing with restricted repacking. *Journal of Combinatorial Optimization*, pages 1–17, 2012.
3. Amotz Bar-Noy, Randeep Bhatia, Joseph Naor, and Baruch Schieber. Minimizing service and operation costs of periodic scheduling. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 11–20, 1998.
4. Amotz Bar-Noy and Richard E Ladner. Windows scheduling problems for broadcast systems. *SIAM Journal on Computing*, 32(4):1091–1113, 2003.
5. Amotz Bar-Noy, Richard E Ladner, and Tami Tamir. Windows scheduling as a restricted version of bin packing. *ACM Transactions on Algorithms (TALG)*, 3(3):28, 2007.
6. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Reallocation problems in scheduling. In *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '13*, pages 271–279, 2013.
7. Joseph Wun-Tat Chan, Tak-Wah Lam, and Prudence WH Wong. Dynamic bin packing of unit fractions items. *Theoretical Computer Science*, 409(3):521–529, 2008.
8. Wun-Tat Chan and Prudence W.H. Wong. On-line windows scheduling of temporary items. In *Proceedings of the 15th International Symposium on Algorithms and Computation*, volume 3341 of *Lecture Notes in Computer Science*, pages 259–270. Springer Berlin Heidelberg, 2005.
9. Xin Han, Chao Peng, Deshi Ye, Dahai Zhang, and Yan Lan. Dynamic bin packing with unit fraction items revisited. *Information Processing Letters*, 110(23):1049–1054, 2010.
10. Mehdi Heydari, Seyed Jafar Sadjadi, and Emran Mohammadi. Minimizing total flow time subject to preemption penalties in online scheduling. *The International Journal of Advanced Manufacturing Technology*, 47(1-4):227–236, 2010.
11. Zoran Ivkovic and Errol L Lloyd. Fully dynamic algorithms for bin packing: Being (mostly) myopic helps. *SIAM Journal on Computing*, 28(2):574–611, 1998.
12. Zhaohui Liu and TC Edwin Cheng. Minimizing total completion time subject to job release dates and preemption penalties. *Journal of Scheduling*, 7(4):313–327, 2004.
13. Peter Sanders, Naveen Sivadasan, and Martin Skutella. Online scheduling with bounded migration. In *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, volume 3142 of *Lecture Notes in Computer Science*, pages 1111–1122. Springer Berlin Heidelberg, 2004.
14. Hadas Shachnai, Tami Tamir, and Gerhard J. Woeginger. Minimizing makespan and preemption costs on a system of uniform machines. *Algorithmica*, 42(3-4):309–334, 2005.
15. SimJava. <http://www.icsa.inf.ed.ac.uk/research/groups/hase/simjava/>, 2006.
16. Jeffery Westbrook. Load balancing for response time. *Journal of Algorithms*, 35(1):1–16, 2000.
17. Haiyan Yu, Yinfeng Xu, and Tengyu Wu. Online inventory replenishment scheduling of temporary orders. *Information Processing Letters*, 113(5):188–192, 2013.

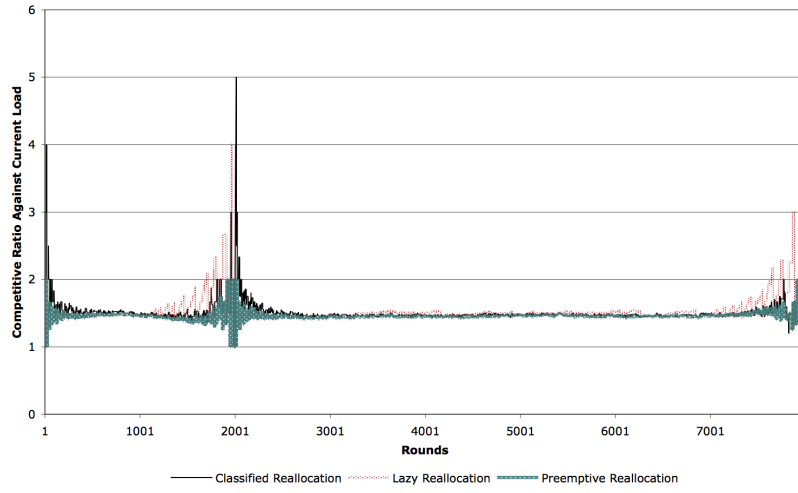


(a) Current load as H_r .

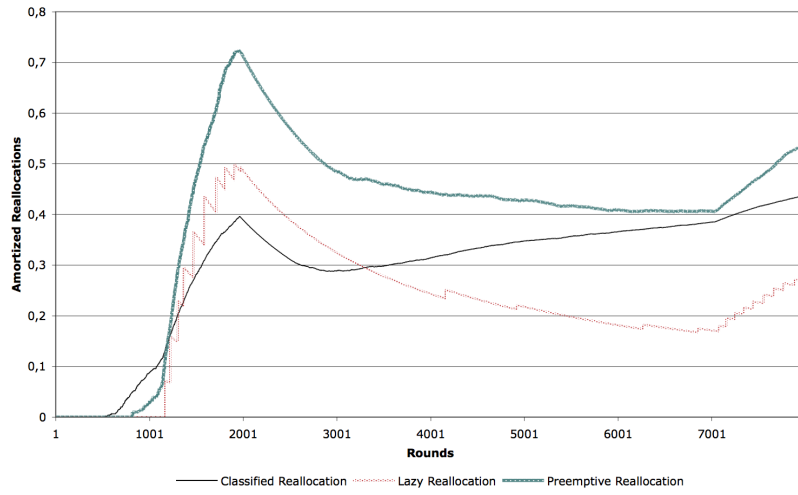


(b) Objective function.

Fig. 1. Performance along rounds for normally distributed inputs.

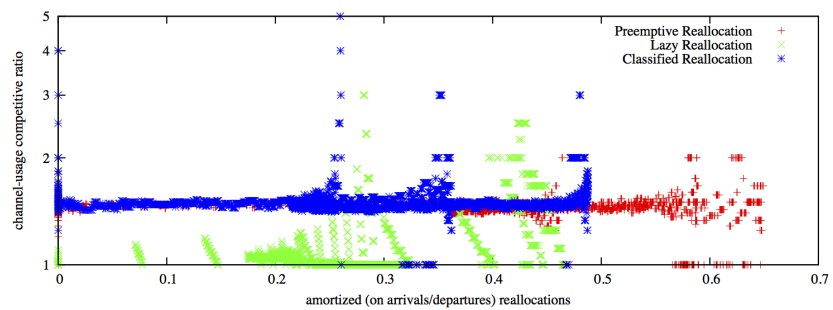


(a) Channel usage competitive ratio.

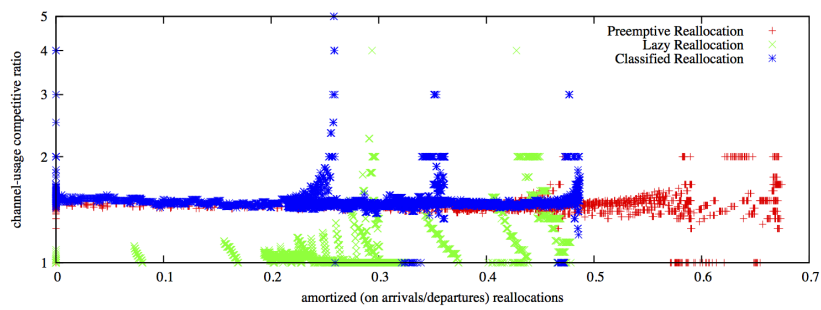


(b) Amortized reallocations.

Fig. 2. Performance along rounds for normally distributed inputs.



(a) Uniform distribution of laxities



(b) Normal distribution of laxities

Fig. 3. Channel usage vs. reallocations